arXiv:1611.09643v1 [math.OC] 29 Nov 2016

Direct solution of piecewise linear systems

Manuel Radons

November 30, 2016

Abstract: *Let S be a real $n \times n$ matrix, $z, \hat{c} \in \mathbb{R}^n$, and $|z|$ the componentwise modulus of z. Then the piecewise linear equation system*

$$z - S|z| = \hat{c}$$

is called an absolute value equation (AVE). It has been proven to be equivalent to the general linear complementarity problem, which means that it is NP hard in general.

We will show that for several system classes the AVE essentially retains the good natured solvability properties of regular linear systems. I.e., it can be solved directly by a slightly modified Gaussian elimination that we call the signed Gaussian elimination. For dense matrices S this algorithm has the same operations count as the classical Gaussian elimination with symmetric pivoting. For tridiagonal systems in n variables its computational cost is roughly that of sorting n floating point numbers. The sharpness of the proposed restrictions on S will be established.

Keywords Absolute value equation; Linear complementarity problem; Piecewise linear equation system; Direct solver; Signed Gaussian elimination

MSC 2010 15A39, 65K05, 90C33

1 Introduction and notation

We denote by $\mathrm{M}_n(\mathbb{R})$ the space of $n \times n$ real matrices, and by $[n]$ the set $\{1, \ldots, n\}$. For vectors and matrices *absolute values and comparisons are used entrywise.* Zero vectors and matrices are denoted by $\mathbf{0}$.

A *signature matrix* Σ, or, briefly, *a signature*, is a diagonal matrix with entries $+1$ or -1. The set of n-dimensional signature matrices is denoted by $\mathrm{diag}_{n,\sigma}$. A single diagonal entry of a signature is a sign σ_i ($i \in [n]$).

Let $S \in \mathrm{M}_n(\mathbb{R})$, $z, \hat{c} \in \mathbb{R}^n$. The *piecewise linear equation system* (PLE)

$$z - S|z| = \hat{c} \tag{1}$$

is called an *absolute value equation* (AVE). It was first introduced by Rohn in [Roh89]. Mangasarian proved its equivalence to the general *linear complementarity problem* (LCP) [MM06]. In [Neu90, pp. 216-230] Neumaier authored a detailed survey about its intimate connection to the research field of *linear interval equations*. A recent result by Griewank and Streubel has shown that PLEs of arbitrary structure can be, with a *one-to-one solution correspondence*, transformed into an AVE [GBRS15, Lem. 6.5].

An especially closely related system type are equilibrium problems of the form

$$Ax + \max(0, x) = b, \tag{2}$$

where $A \in \mathrm{M}_n(\mathbb{R})$ and $x, b \in \mathbb{R}^n$. (A prominent example is the first hydrodynamic model presented in [BC08].) Using the identity $\max(s, t) = (s + t + |s - t|)/2$, equality (2) can be reformulated as

$$Ax + \frac{x + |x|}{2} = b \quad \Longleftrightarrow \quad (2A + I)x + |x| \equiv Bx + |x| = 2b. \tag{3}$$

For regular B, system (3) is clearly equivalent to (1).

This position at the crossroads of several interesting problem areas gives relevance to the task of developing efficient solvers for the AVE. The latest publications on the matter include approaches by linear programming [Man14] and concave minimization [Man07a], as well as a variety of Newton and fixed point methods (see, e.g., [BC08], [YY12], [HHZ11] or [GBRS15]).

Let $\Sigma \in \mathrm{diag}_{n,\sigma}$ s.t. $\Sigma z = |z|$. (Note that, since $0 = +0 = -0$, we need no "0"-sign.) Then we can rewrite (1) as

$$(I - S\Sigma)z = \hat{c}. \tag{4}$$

In this form it becomes apparent that the main difficulty in the computation of a solution for (4) is to determine the proper signature Σ for z. That is, to determine in which of the 2^n orthants about the origin z lies. This is NP-hard in general [Man07b].

It was proven by Rump in [Rum97, Cor. 2.9] that checking the system for unique solvability is NP-hard as well, as it is equivalent to checking whether a quantity called the *sign-real spectral radius* of S is smaller than one, which in

turn is equivalent to checking whether the system matrix of the equivalent LCP is a *P-matrix*. As these notions and results are fundamental to the understanding of the AVE, we will give a short account of them in the second section. There we will also see that the systems investigated in the present paper, for all of which it holds $\|S\|_\infty < 1$, are uniquely solvable.

The following simple observation is key to the subsequent discussion:

Proposition 1.1. *Let $S \in \mathrm{M}_n(\mathbb{R})$ and $z, \hat{c} \in \mathbb{R}^n$ such that they satisfy (4). Then, if $\|S\|_\infty < 1$, for at least one $i \in [n]$ the signs of z_i and $\hat{c}_i$ have to coincide.*

Proof. Let z_i be an entry of z s.t. $|z_i| \geq |z_j|$ for all $j \in [n]$. If $z_i = 0$, then $z = \mathbf{0}$ and thus $\hat{c} \equiv z - S|z|$ is the zero vector as well – and the statement holds trivially. If $|z_i| > 0$, then $\left|e_i^T S|z|\right| < |z_i|$, due to the norm constraint on S. Thus, $\hat{c}_i = z_i - e_i^T S|z|$ will adopt the sign of z_i. □

We do not know though, for which indices the signs coincide. In the third section we will derive several types of structural restrictions on S, each of which will guarantee the coincidence of the signs of z_i and $\hat{c}_i$ for all $i \in [n]$ with $|\hat{c}_i|$ maximal in $\hat{c}$.

In the fourth paragraph we will devise a modified Gaussian elimination that exploits this knowledge. This *signed Gaussian elimination* (SGE) will base on the following central points:

- We are enabled to perform one step of Gaussian elimination on the AVE in the form (4), if we know the correct sign of z_1.
- If $\|S\|_\infty < 1$, no row or column pivot causes numerical instabilities in the performance of a Gaussian elimination step on system (4). Hence, we can always produce a constellation, where $|\hat{c}_1|$ is maximal in $\hat{c}$.
- The restrictions on S developed in the third paragraph are invariant under Gaussian elimination steps.

For S that conform to the restrictions derived in paragraph three, the first two points mean that we can always perform one Gaussian elimination step on system (4). The third point ensures that we can repeat the procedure for the reduced system(s) and ultimately calculate the correct (unique) solution of the AVE.

We will briefly analyze the modified algorithm's runtime in the dense and tridiagonal case. For a dense matrix S the SGE has the operations count of a Gaussian elimination with symmetric pivoting. For the tridiagonal SGE the supplementary operations cost roughly as much as sorting $\hat{c}$ with respect to the absolute value of its entries. As the underlying tridiagonal Gaussian elimination,

also known as the Thomas Algorithm, is in $\mathcal{O}(n)$, this means that the asymptotical complexity of the modified algorithm depends, at the current state of research, one-to-one on the implementation of the extra effort.

The paper is concluded by a discussion of the sharpness of the proposed restrictions on S.

For readers primarily interested in the algorithmic results, we remark that inequality (6), equivalence 1. $\Leftrightarrow$ 3. from Theorem 2.1., and the statements of Theorem 3.1. present the most basic preknowledge that should enable them to work with the fourth paragraph.

Note that we already outlined the approach described above in [GBRS15, Parag. 7]. This paper presents the announced elaboration on the concept.

2 Sign-real spectral radius

Denote by $\rho(S)$ the spectral radius of S and let

$$\rho_0(S) \equiv \max\{|\lambda| : \lambda \textit{ real} \text{ eigenvalue of } S\}$$

be the *real spectral radius* of S. Then its sign-real spectral radius is defined as follows (see [Rum97, Def. 1.1]):

$$\rho_0^s(S) \;\equiv\; \max\left\{\rho_0(\Sigma S) : \Sigma \in \mathrm{diag}_{n,\sigma}\right\}.$$

The exponential number of signatures Σ accounts for the NP-hardness of the computation of $\rho_0^s(S)$. It is easy to check that $\mathrm{diag}_{n,\sigma}$ is a finite subgroup of $\mathrm{Gl}_n(\mathbb{R})$. Thus, for a fixed signature $\bar{\Sigma}$, the sets $\{\Sigma(\bar{\Sigma}S) : \Sigma \in \mathrm{diag}_{n,\sigma}\}$ and $\{\Sigma S : \Sigma \in \mathrm{diag}_{n,\sigma}\}$ are identical modulo a permutation. Furthermore, since all $\Sigma \in \mathrm{diag}_{n,\sigma}$ are obviously involutive, i.e., $\Sigma^{-1} = \Sigma$, the spectra of S and $\Sigma S \Sigma$ are identical. These observations immediately yield the useful identity

$$\rho_0^s(S) = \rho_0^s(\Sigma_1 S) = \rho_0^s(S\Sigma_2) = \rho_0^s(\Sigma_1 S \Sigma_2) \qquad \forall\ \Sigma_1, \Sigma_2 \in \mathrm{diag}_{n,\sigma}\,.$$

Recall that a real (or complex) square matrix is called a P-matrix if every principal minor is positive [CPS92, p. 147]. An LCP has a unique solution for all right hand sides if and only if its system matrix is a P-matrix [CPS92, p. 148, Thm. 3.3.7]. We will now (re-) prove some essential facts about the relation between $\rho_0^s(S)$ and the solvability properties of (4).

Theorem 2.1. *Let $S \in \mathrm{M}_n(\mathbb{R})$. Then the following are equivalent:*

1. $\rho_0^s(S) \;<\; 1$.

2. $(I - S)^{-1}(I + S)$ *is a P-matrix.*

3. *The system* $(I - S\Sigma)z = \hat{c}$ *has a unique solution for all* $\hat{c} \in \mathbb{R}^n$.

4. *The function* $\varphi : \mathbb{R}^n \to \mathbb{R}^n,\ z \to z + S|z|$ *is bijective.*

5. *For all* $\hat{c} \in \mathbb{R}^n$ *there exists a unique* $\Sigma \in \mathrm{diag}_{n,\sigma}$ *s.t.* $b \equiv (I - S\Sigma)^{-1}\hat{c}$ *lies in the orthant defined by* Σ.

6. $\det(I - S\Sigma) > 0$ *for all* $\Sigma \in \mathrm{diag}_{n,\sigma}$.

7. $\det(I - SD) > 0$ *for all real diagonal matrices* $D \in \mathrm{M}_n(\mathbb{R})$ *with* $\|D\|_\infty \leq 1$.

Proof. Let $\mu, \lambda \in \mathbb{R}$. We note that, since $\det((A + \mu I) - \lambda I) = \det(A - (\lambda - \mu)I)$, the spectrum of a matrix $\bar{A} \equiv A + \mu I$ is the spectrum of A shifted by μ along the real axis. We will refer to this fact by the abbreviation SHIFT.

1. $\Rightarrow$ 6. : Let $\rho_0^s(S) < 1$ and fix a signature Σ. Then the absolute value of all real eigenvalues of $S\Sigma$ is smaller than one. Thus, by SHIFT, all real eigenvalues of $(I - S\Sigma)$ lie in the open interval $(0, 2)$, which means that their product is positive. The complex eigenvalues appear in conjugate pairs, hence their product is positive as well. This yields a positive determinant.

6. $\Rightarrow$ 7. : Let $\Sigma, \Sigma' \in \mathrm{diag}_{n,\sigma}$ be signatures that differ only in the first sign. By assumption we have

$$\det(I - S\Sigma) > 0 \text{ and } \det(I - S\Sigma') > 0 \ .$$

Then, by the linearity of the determinant for rank-1 updates, it holds

$$\det(I - SD) > 0 \ ,$$

where $D \in \mathrm{M}_n(\mathbb{R})$ is a diagonal matrix whose first entry lies in the interval $[-1, 1]$, while all others equal the corresponding entries in Σ / Σ'. Now apply this argument inductively.

7. $\Rightarrow$ 1. : Assume that $\det(I - S\Sigma) > 0$ for all for all real diagonal matrices $D \in \mathrm{M}_n(\mathbb{R})$ with $\|D\|_\infty \leq 1$, but $\rho_0^s(S) \geq 1$. Then there exists a signature $\Sigma \in \mathrm{diag}_{n,\sigma}$ s.t. $S\Sigma$ has at least one real eigenvalue λ with $|\lambda| \geq 1$. Define $D := \frac{1}{\lambda}\Sigma$. Clearly, D is a diagonal matrix with $\|D\|_\infty \leq 1$. And, by SHIFT, it holds

$$\det(I - SD) = 0$$

– in contradiction to the hypothesis.

2. $\Leftrightarrow$ 3. : Let $z \equiv u - w$ with $u \perp w$ in that $u \geq 0 \leq w$ and $u^\top w = 0$, we obtain $|z| = u + w$. Substituting this into the AVE, we get

$$\begin{aligned} & \hat{c} &=& \; u - w + S(u + w) \\ \iff \qquad & (I - S)w &=& -\hat{c} + (I + S)u \\ \iff \qquad & w &=& -(I - S)^{-1}\hat{c} + (I - S)^{-1}(I + S)u. \end{aligned}$$

The latter equation has the form of an LCP and hence possesses a unique solution if and only if $(I - S)^{-1}(I + S)$ is a P-matrix.

3. $\Leftrightarrow$ 5. : If b lies in the orthant defined by Σ, then $\Sigma b = |b|$, that is, b is a solution of the system. But then the equivalence is clear.

2. $\Leftrightarrow$ 7. : For $A, B \in \mathrm{M}_n(\mathbb{R})$ the following equivalency holds: $TA + (I - T)B$ is regular for all n-dimensional diagonal matrices T with entries $t_i \in [0, 1] \iff A^{-1}B$ is a P-matrix [JT95, Thm. 3.4].

But we have $T(I-S)+(I-T)(I+S) = I-(I-2T)S$ – and the set of matrices $I - 2T$ is clearly identical to the set of diagonal matrices D with $\|D\|_\infty \leq 1$.

7. $\Rightarrow$ 6. : Obvious, since the n-dimensional diagonal matrices D with $\|D\|_\infty \leq 1$ are the convex hull of $\mathrm{diag}_{n,\sigma}$.

6. $\Rightarrow$ 4. : If we interpret (1) as the piecewise linear function

$$\varphi : \mathbb{R}^n \;\to\; \mathbb{R}^n,\; z \;\to\; z + S|z| \;,$$

then $\det(I - S\Sigma) > 0$ for all signatures means that the limiting Jacobians of φ all have the same determinant sign – a property which is called coherent orientation and implies surjectivity of the map [Sch12, p. 32]. Since the piecewise linearity of φ originates in absolute values that are not encapsulated in other absolute values, it is a simply switched piecewise linear function in the sense of [GBRS15, Parag. 2].

Also, by continuity of the determinant, there exists, for each signature Σ, an open neighborhood $M_\Sigma \subset \mathrm{M}_n(\mathbb{R})$ about I s.t. for all $\bar{I} \in M_\Sigma$ we have $\det(\bar{I} - S\Sigma) > 0$. Then (the finite intersection of open sets) $M \equiv \bigcap_{\Sigma \in \mathrm{diag}_{n,\sigma}} M_\Sigma$ is a nonempty open neighborhood about I s.t. $\det(\bar{I} - S\Sigma) > 0$ for all $\bar{I} \in M$ and all $\Sigma \in \mathrm{diag}_{n,\sigma}$. Hence, the coherent orientation of φ is stable under small perturbations of I. Thus, it conforms to the definition of a stably coherently oriented and simply switched piecewise linear map in [GBRS15, Parag. 4.]. And as such it is also injective (see [GBRS15, Cor. 4.5.]), hence bijective.

4. $\Rightarrow$ 3. : Obvious. □

Remark 2.1. *The equivalency* 1. $\Leftrightarrow$ 3. $\Leftrightarrow$ 6. $\Leftrightarrow$ 7. *in Theorem 2.1. was first stated by J. Rohn. The new proofs (mostly) use linear complementarity theory*

and thus showcase the kinship of LCPs and AVEs. For the linear algebraic original proofs, see, e.g., [Neu90, p. 218]. In [GBRS15, Parag. 7.5] Griewank proposed the transformation of a general PLE in so-called abs-normal representation into an LCP. To prove 2. $\Leftrightarrow$ 3. *we adapted this reformulation to the AVE. The proof of* 6. $\Rightarrow$ 4. *demonstrates the productive capacity of recent piecewise linear theory. The fact that the linear transformation* $(I - S\Sigma)^{-1}$ *maps* $\hat{c}$ *to a different orthant than the one defined by* Σ *for all but one* Σ *in* $\mathrm{diag}_{n,\sigma}$ *– that is, point* 5. *– is not interesting in the present setting, but gains significance in the context of Newton type approaches to the solution of* (4)*, such as those presented in [GBRS15].*

Note that, by the SHIFT argument in the proof, the above statements still hold if we replace I by $\alpha \cdot I$ and $\rho_0^s(S) < 1$ by $\rho_0^s(S) < \alpha$, respectively (α a positive scalar).

Furthermore, if we keep in mind that multiplication by a signature matrix merely flips the signs of a row or column without changing the absolute values of the entries, we immediately see:

$$\|\Sigma_1 S\|_\infty = \|S\Sigma_2\|_\infty = \|\Sigma_1 S\Sigma_2\|_\infty = \|S\|_\infty \ \forall\ \Sigma_1, \Sigma_2 \in \mathrm{diag}_{n,\sigma}\,. \tag{5}$$

Consequently, we also get:

$$\rho_0(\Sigma S) \leq \rho(\Sigma S) \leq \|\Sigma S\|_\infty = \|S\|_\infty \ \forall\ \Sigma \in \mathrm{diag}_{n,\sigma}, \tag{6}$$

which implies $\rho_0^s(S) \leq \|S\|_\infty$. As we we will only consider S with $\|S\|_\infty < 1$ in the present work, this yields $\rho_0^s(S) < 1$ for all systems investigated hereafter and thus positively answers the question of their unique solvability.

While we only make use of the infinity-case, it is worth mentioning that $\rho_0^s(S)$ is, in fact, bounded by all p-norms (see [Rum97, Thm. 2.15]). Moreover, note that by the Perron-Frobenius rescaling introduced in [GBRS15, Lem. 6.4] any system $(I - S\Sigma)z = \hat{c}$ with $\|S\|_1 < 1$ can be transformed into a system $(I - S'\Sigma)z' = c'$ with $\|S'\|_\infty < 1$.

3 Main theorem

We continue to use $S \in \mathrm{M}_n(\mathbb{R})$ and $z, \hat{c} \in \mathbb{R}^n$ in their roles of the previous sections, and introduce a slight abuse of notation: Hereafter we will identify a vector $v \in \mathbb{R}^n$ with the set of its entries, ordered by their index. That is, $v \equiv \{v_1, \ldots, v_n\}$. This way, the sets $C_{\max}$ and $\Sigma_{\neq}$ in the definitions below can contain arbitrary subsets of the entries of $\hat{c}$. Now let

$$C_{\max} \ \equiv\ \{\hat{c}_j \in \hat{c} : |\hat{c}_j| = \max_{k \in [n]}(|\hat{c}_k|)\},$$

and

$$\Sigma_{\neq} \;\equiv\; \{\hat{c}_j \in \hat{c} : \mathbf{sign}(z_j) \neq \mathbf{sign}(\hat{c}_j)\},$$

where **sign** denotes the *signum function*. That is, **sign** is an element in {-1, 0, 1}. This is a stricter notion of sign coincidence than the one given in the introduction, where 0 was essentially treated as a logical *don't-care*, for which both + and − were allowed as proper signs.

Theorem 3.1. *Let $S \in \mathrm{M}_n(\mathbb{R})$ and $z, \hat{c} \in \mathbb{R}^n$ such that it holds*

$$(I - S\Sigma)z \;=\; \hat{c}\ ,$$

where $\Sigma z = |z|$. Then we have

$$C_{\max} \cap \Sigma_{\neq} \;=\; \emptyset\ ,$$

if one of the following conditions is satisfied:

1. $\|S\|_\infty < \frac{1}{2}$.
2. *S is irreducible with* $\|S\|_\infty \leq \frac{1}{2}$.
3. *S is strictly diagonally dominant with* $\|S\|_\infty \leq \frac{2}{3}$.
4. *S is tridiagonal with* $\|S\|_\infty < 1$.

3.1 Proof of 1. and 2.

The following lemma will provide a sufficient condition for the statement of the theorem to hold.

Lemma 3.1. *Let $\hat{c}, z \in \mathbb{R}^n$, $S \in \mathrm{M}_n(\mathbb{R})$ with $\rho_0^s(S) < 1$, and $\Sigma \in \mathrm{diag}_{n,\sigma}$, such that they satisfy* (4)*. Then, if the matrix $A \equiv (I - S\Sigma)^{-1}$ is strictly diagonally dominant with a positive diagonal, we have*

$$C_{\max} \cap \Sigma_{\neq} \;=\; \emptyset.$$

Proof. Fix any $\hat{c}_i \in C_{\max}$. We distinguish two cases:

Case 1: Let $|\hat{c}_i| > 0$. Since $|\hat{c}_i| \geq |\hat{c}_j|$ for all $j \in [n]$, we always have $|\hat{c}_i a_{ii}| > \sum_{j \neq i} |\hat{c}_j a_{ij}|$ due to the strict diagonal dominance of A. Consequently, since a_{ii} is positive, z_i will adopt the sign of $\hat{c}_i$.

Case 2: Let $\hat{c}_i = 0$. Then $\hat{c} = \mathbf{0}$, as $\hat{c}_i \in C_{\max}$. Hence, because of the unique solvability implied by $\rho_0^s(S) \leq |S|_\infty < 1$, z is the zero vector as well – which especially means $z_i = 0$. □

With this criterium in hand, we can prove the first two statements of the theorem:

Lemma 3.2. *Let $S \in \mathrm{M}_n(\mathbb{R})$ be irreducible with $\|S\|_\infty \leq \frac{1}{2}$, then the inverse of the matrix $A \equiv I - S$ is strictly diagonally dominant and has a positive diagonal.*

Proof. We have $\|S^k\|_\infty \leq \|S\|_\infty^k \leq \frac{1}{2^k}$, which implies $\lim_{k\to\infty}(I-A)^k = \lim_{k\to\infty} S^k = 0$. Thus, A^{-1} can be expressed via the Neumann series

$$A^{-1} = \sum_{k=0}^{\infty}(I-A)^k = \sum_{k=0}^{\infty} S^k = I + \sum_{k=1}^{\infty} S^k.$$

The inequality $\| \sum_{k=1}^{\infty} S^k\|_\infty \leq \sum_{k=1}^{\infty} \|S\|_\infty^k \leq \sum_{k=1}^{\infty} \frac{1}{2^k} = 1$ already ensures weak diagonal dominance of A^{-1}.

Now fix any $i \in [n]$ and assume that the i-th row of A^{-1} were not strictly dominated by its diagonal entry. Denote the entries of S^k by $s_{ij}^{(k)}$ for $i, j \in [n]$. Then

$$1 = \sum_{j=1}^{n} \left| \sum_{k=1}^{\infty} s_{ij}^{(k)} \right| \leq \sum_{j=1}^{n} \sum_{k=1}^{\infty} |s_{ij}^{(k)}| \leq \sum_{k=1}^{\infty} \frac{1}{2^k} = 1,$$

which implies that

$$\left| \sum_{k=1}^{\infty} s_{ij}^{(k)} \right| = \sum_{k=1}^{\infty} |s_{ij}^{(k)}| \quad \forall j \in [n] \tag{7}$$

and $\sum_{j=1}^{n} |s_{ij}^{(k)}| = \frac{1}{2^k}$ for all $k \geq 1$. In particular,

$$\begin{aligned} \frac{1}{2^{k+1}} &= \sum_{j=1}^{n} |s_{ij}^{(k+1)}| = \sum_{j=1}^{n} \left| \sum_{i=1}^{n} s_{ir}^{(k)} s_{rj}^{(1)} \right| \leq \sum_{j=1}^{n} \sum_{r=1}^{n} |s_{ir}^{(k)} s_{rj}^{(1)}| \\ &= \sum_{r=1}^{n} |s_{ir}^{(k)}| \sum_{j=1}^{n} |s_{rj}^{(1)}| \leq \sum_{r=1}^{n} |s_{ir}^{(k)}| \cdot \frac{1}{2} = \frac{1}{2^{k+1}}, \end{aligned}$$

which implies for each $k \geq 1$ that

$$\left| \sum_{r=1}^{n} s_{ir}^{(k)} s_{rj}^{(1)} \right| = \sum_{r=1}^{n} |s_{ir}^{(k)} s_{rj}^{(1)}| \quad \forall j \in [n]. \tag{8}$$

Claim: For each $k \geq 1$, the i-th row of S^k has the same entry pattern as the i-th row of $|S|^k$. We prove this by induction. The case $k = 1$ is trivial. Assume the

claim holds for a given k. Let $\mathcal{I}_i^{(k)} = \{a_1, \ldots, a_m\}$ be the set of indices of the nonzero entries of the i-th row S^k, or equivalently of $|S|^k$. Define $\mathcal{I}_{a_1}^{(1)}, \ldots, \mathcal{I}_{a_m}^{(1)}$ analogously, and let $\mathcal{I} \equiv \bigcup_{a \in \mathcal{I}_i^{(k)}} \mathcal{I}_a^{(1)}$. Obviously, $\mathcal{I}$ is precisely the set of indices of the nonzero entries in the i-th row of $|S|^{k+1} = |S|^k|S|$, and $s_{ij}^{(k+1)} \neq 0$ at most if $j \in \mathcal{I}$. But this necessary condition is also sufficient because otherwise (8) would be violated. This completes the proof of the claim.

Since $|S|$ is irreducible and nonnegative, there exists a power $|S|^{k_i}$ with a positive entry at (i,i) (see, e.g., [Kit98, p. 3]). By what we just showed, this implies $s_{ii}^{(k_i)} \neq 0$. Therefore $(s_{ii}^{(k_i)})^2 > 0$ and by (8) also $s_{ii}^{(2k_i)} > 0$. Now (7) implies that $s_{ii}^{(k)} \geq 0$ for all $k \geq 1$.

Let D be the diagonal part of $\sum_{k=1}^{\infty} S^k$ and $B \equiv \sum_{k=1}^{\infty} S^k - D$. Then $A^{-1} = I + D + B$, where $(I+D)_{ii} \geq 1 + s_{ii}^{(2k_i)} > 1$, while $\sum_{j=1}^{n} |B_{ij}| \leq 1 - s_{ii}^{(2k_i)} < 1$. So our assumption that the i-th row of A^{-1} were not strictly dominated by its diagonal entry is in fact wrong. This completes the proof. □

Note that for $\|S\|_1 \leq \frac{1}{2}$ the arguments of the proof imply strict diagonal dominance of the inverse over the columns. Obviously, we also have strict diagonal dominance of $(\alpha A)^{-1} = \frac{1}{\alpha}(A^{-1})$, where $\alpha \in \mathbb{R}\backslash\{0\}$.

Corollary 3.3. *Let $S \in \mathrm{M}_n(\mathbb{R})$ with $\|S\|_\infty < \frac{1}{2}$, then the inverse of the matrix $A \equiv I - S$ is strictly diagonally dominant and has a positive diagonal.*

Proof. Consider the Neumann series $\sum_{k=0}^{\infty} S^k = I + \sum_{k=1}^{\infty} S^k$ in the proof above. With the sharper bound we get $\| \sum_{k=1}^{\infty} S^k \|_\infty \leq \sum_{k=1}^{\infty} \|S\|_\infty^k < \sum_{k=1}^{\infty} \frac{1}{2^k} = 1$. This ensures that A^{-1} is strictly diagonally dominant with a positive diagonal. □

Now recall that, by (5), we have $\|S\Sigma\|_\infty = \|S\|_\infty$ for all $\Sigma \in \mathrm{diag}_{n,\sigma}$. Then it is clear that the restrictions stated in 1. and 2. in Theorem 3.1. imply the strict diagonal dominance of $(I - S\Sigma)^{-1}$ for all $\Sigma \in \mathrm{diag}_{n,\sigma}$ – which also includes the proper signature of the solution (in the sense that $\Sigma z = |z|$) and thus allows for the application of Lemma 3.1. to the situation of the first two conditions. This completes the proof of the first two statements of Theorem 3.1.

Remark 3.4. *The matrix*

$$S \equiv \begin{bmatrix} 0 & \frac{1}{2} \\ 0 & \frac{1}{2} \end{bmatrix} \qquad \textit{with} \qquad (I-S)^{-1} = \begin{bmatrix} 1 & 1 \\ 0 & 2 \end{bmatrix}$$

shows that in the limiting case $\|S\|_\infty = \frac{1}{2}$ the criterium of irreducibility cannot be omitted in Lemma 3.2. Furthermore, for $\epsilon > 0$ arbitrarily small

$$S \equiv \begin{bmatrix} \epsilon & \frac{1}{2} \\ 0 & \frac{1}{2} \end{bmatrix} \qquad \textit{with} \qquad (I-S)^{-1} = \begin{bmatrix} 1-2\epsilon & 1+2\epsilon \\ 0 & 2-2\epsilon \end{bmatrix}$$

and $\|S\|_\infty = \frac{1}{2} + \epsilon$ *proves the sharpness of the bound* $\|S\|_\infty \leq \frac{1}{2}$.

Also note that, if $S \in \mathrm{M}_n(\mathbb{R})$ *is nilpotent (which implies the nilpotency of* $S\Sigma$ *for all* $\Sigma \in \mathrm{diag}_{n,\sigma}$*), the Neumann expansion of* $(I - S\Sigma)^{-1}$ *has at most* n *summands. Thus, if* $\|S\|_\infty \leq \frac{1}{2}$*, we have*

$$\left\| \sum_{k=1}^{\infty} S^k \right\|_\infty \;\leq\; \sum_{k=1}^{\infty} \|S\|_\infty^k \;=\; \sum_{k=1}^{n-1} \|S\|_\infty^k \;\leq\; \sum_{k=1}^{n-1} \frac{1}{2^k} \;<\; 1$$

– and again obtain strict diagonal dominance of the inverse of $I - S\Sigma$*. But, since nilpotent matrices are permutationally similar to strictly upper triangular matrices, the corresponding AVEs can be solved by a modified backwards substitution in* $\mathcal{O}(n^2)$ *operations. Which is why we did not include this case in the main theorem.*

3.2 Proof of 3.

Denote by $\mathrm{diag}_n(a_1, \ldots, a_n)$ the n-dimensional diagonal matrix with entries $a_1, \ldots, a_n \in \mathbb{R}$. Furthermore, define analogously to $C_{\max}$:

$$Z_{\max} \equiv \{z_j \in z : |z_j| = \max_{k \in [n]}(|z_k|)\} \;,$$
$$Z_{\min} \equiv \{z_j \in z : |z_j| = \min_{k \in [n]}(|z_k|)\} \;.$$

We first exclude two special cases:

- As $\rho_0^s(S) \leq \|S\|_\infty < 1$, the system is uniquely solvable and the statement thus holds trivially for $z = \mathbf{0}$. We therefore limit our attention to cases, where z has at least one nonzero entry.

- $Z_{\min}$ and $Z_{\max}$ are either disjoint or equal. In both cases neither set is empty. Since $\|S\|_\infty < 1$, it is $\mathbf{sign}(z_i) = \mathbf{sign}(\hat{c}_i)$ for all $i \in [n]$, if $|z_1| = \cdots = |z_n|$, i.e., if $Z_{max} = Z_{\min}$. Thus, we only have to prove 3. for cases, where $Z_{\max} \neq Z_{\min}$ and hence both sets are disjoint.

The following observation is crucial:

- If $\|S\|_\infty < 1$ and $z_i \in Z_{\max}$, we have $\sum_j |s_{ij} z_j| < |z_i|$ and hence $\mathbf{sign}(\hat{c}_i) = \mathbf{sign}(z_i)$. Consequently, if there were a tuple $(S, z, \hat{c})$ that violated the claim of the theorem, for any $\hat{c}_j \in C_{\max} \cap \Sigma_{\neq}$ we would have $z_j \notin Z_{\max}$.

The proof is performed by induction. For $n = 1$ the statement holds trivially. Assume it holds for $N \geq 1$, but there exists a tuple $(S, z, \hat{c})$ in dimension $N + 1$ that falsifies it. We distinguish two cases:

Case 1: Let $\hat{c}_i \in C_{\max}$ and $z_i \notin Z_{\min}$ s.t. $\mathbf{sign}(\hat{c}_i) \neq \mathbf{sign}(z_i)$. We will, from the falsifying tuple $(S, z, \hat{c})$ in dimension $N+1$, construct a tuple $(\bar{S}, \bar{z}, \bar{c})$ in dimension N that falsifies the statement as well and thus contradicts the induction hypothesis.

Assume *w.l.o.g.* that $z_{N+1} \in Z_{\min}$. Then for all $j \in [N]$ there exists a scalar $\zeta_j \in [0, 1]$ such that

$$\zeta_j \cdot |z_j| = |z_{N+1}| \quad \Longrightarrow \quad \zeta_j \cdot s_{j,N+1} \cdot |z_j| = s_{j,N+1} \cdot |z_{N+1}| \ .$$

Denote by $S_{N+1,N+1}$ an N-dimensional square matrix derived from S by removing row and column $N+1$. Then we have, for

$$\bar{z} \ = \ (z_1, \ldots, z_N)^T \tag{9}$$

and

$$\bar{S} \ \equiv \ S_{N+1,N+1} \ + \ \mathrm{diag}_N(\zeta_1 \cdot s_{j,N+1}, \ldots, \zeta_N \cdot s_{N,N+1}) \tag{10}$$

that

$$\bar{z} + \bar{S}|\bar{z}| \ = \ (\hat{c}_1, \ldots, \hat{c}_N)^T \equiv \ ^{-}c \ . \tag{11}$$

Since the coefficients ζ_i are in $[0, 1]$ for all $i \in [n]$, we have $\|\bar{S}\|_\infty \leq \|S\|_\infty \leq \frac{2}{3}$. For the same reason $\bar{S}$ is also still strictly diagonally dominant. Now, since $z_i \notin Z_{\min}$, but $z_{N+1} \in Z_{\min}$, we must have $1 \leq i \leq N$. That is, row i (that contains the contradiction) was not removed by the construction. Thus, the tuple $(\bar{S}, \bar{z}, \bar{c})$ contradicts the induction hypothesis for dimension N.

Case 2: Let $\hat{c}_i \in C_{\max}$ and $z_i \in Z_{\min}$ s.t.

$$\mathbf{sign}(\hat{c}_i) \ \neq \ \mathbf{sign}(z_i) \ .$$

There is the possibility that z_i is the only element in $Z_{\min}$. In this case the construction devised above fails, as it eliminates the row that contains the contradiction. We thus use an approach by direct computation. For this we note that, since $\|S\|_\infty < 1$ and thus $s_{ii} < 1$ for all $i \in [n]$, the following two statements hold:

$$\mathbf{sign}(z_i - s_{ii}|z_i|) \ = \ \mathbf{sign}(z_i)$$

and

$$|z_i - s_{ii}|z_i|| \ \geq \ (1 - |s_{ii}|)|z_i| \ . \tag{12}$$

With $\hat{c}_i = z_i - \sum_{j=1}^{N+1} s_{ij}|z_j|$, and since $\mathbf{sign}(\hat{c}_i) \neq \mathbf{sign}(z_i) = \mathbf{sign}(z_i - s_{ii}|z_i|)$, it holds

$$\mathbf{sign}(z_i - s_{ii}|z_i|) \;\neq\; \mathbf{sign}\left(-\sum_{j\neq i} s_{ij}|z_j|\right)$$

and thus

$$|\hat{c}_i| \;=\; \left||z_i - s_{ii}|z_i|| - \left|\sum_{j\neq i} s_{ij}|z_j|\right|\right| \;.$$

Using (12) then yields:

$$|\hat{c}_i| \;=\; \left|z_i - s_{ii}|z_i| - \sum_{j\neq i} s_{ij}|z_j|\right| \tag{13}$$

$$\leq\; \left|(1-|s_{ii}|)|z_i| - \sum_{j\neq i}|s_{ij}z_j|\right| \;\leq\; \sum_{j\neq i}|s_{ij}z_j| \;. \tag{14}$$

Furthermore, from $\sum_j |s_{ij}| \leq \frac{2}{3}$ (norm constraint) and $\sum_{j\neq i}|s_{ij}| < s_{ii}$ (strict diagonal dominance), we get $\sum_{j\neq i}|s_{ij}| < \frac{1}{3}$. Now let $z_m \in Z_{\max}$. With (13) and (14) we get the leftmost inequality in:

$$|\hat{c}_i| \;\leq\; \sum_{j\neq i}|s_{ij}z_j| \;\leq\; \sum_{j\neq i}|s_{ij}z_m| \;=\; |z_m|\cdot\sum_{j\neq i}|s_{ij}| \;<\; \frac{1}{3}|z_m| \;. \tag{15}$$

But we also have:

$$|\hat{c}_m| \;=\; \left|z_m - \sum_j s_{mj}z_j\right| \;\geq\; \left||z_m| - \sum_j |s_{mj}z_m|\right| \;\geq\; \frac{1}{3}|z_m| \;.$$

Together with (15) the latter inequality gives $|\hat{c}_m| > |\hat{c}_i|$ – which contradicts $\hat{c}_i \in C_{\max}$ and completes the proof of Theorem 3.1.3.

3.3 Proof of 4.

The proof is again inductive. The case $n = 2$ follows from a straightforward elementary calculation for which we refer to the appendix of [Rad16]. Now assume the statement of the theorem would hold for an $N \geq 2$, but the tuple $(S, z, \hat{c})$

would contradict it in dimension $N+1$. We duplicate the argument from *Case 1* in the proof of *3.*:

As $N+1 \geq 3$, we can organize the system *w.l.o.g.* such that $z_N \in Z_{\max}$ and $\hat{c}_{N+1} \notin C_{\max}$. Then there exists a scalar $\zeta \in [0,1]$ such that

$$\zeta \cdot |z_N| = |z_{N+1}| \quad \Longrightarrow \quad \zeta \cdot s_{j,N+1} \cdot |z_N| = s_{j,N+1} \cdot |z_{N+1}|.$$

Then

$$\bar{S} \;\equiv\; S_{N+1,N+1} + \mathrm{diag}_N(0,\ldots,0,\zeta s_{N,N+1})$$

is still symmetric and tridiagonal with $\|S\|_\infty < 1$. And, for $\bar{z}, \bar{c}$ defined as in (9) and (11), respectively, we have $\bar{z} + \bar{S}|\bar{z}| = \bar{c}$. Thus, the tuple $(\bar{S}, \bar{z}, \bar{c})$ contradicts the induction hypothesis for dimension N.

This completes the proof of the fourth statement and thus of the main theorem.

4 Signed Gaussian elimination

4.1 Preliminaries

We will show the three bullet points stated in the introduction: Let S and Σ as in (4) and define the following matrix-block partitions:

$$\Sigma \equiv \begin{bmatrix} \sigma_1 & 0 \\ 0 & \bar{\Sigma} \end{bmatrix} \qquad \text{and} \qquad S \equiv \begin{bmatrix} E & F \\ G & H \end{bmatrix}, \tag{16}$$

where $\sigma_1 \in \{+1,-1\}$ is the first diagonal entry – i.e., the first sign – of Σ and $E \equiv s_{11}$. Then the first step of a Gaussian elimination will transform $(I - S\Sigma)$ into

$$\begin{bmatrix} 1-\sigma_1 s_{11} & -F\bar{\Sigma} \\ 0 & I - H\bar{\Sigma} + \sigma_1 G(1-\sigma_1 s_{11})^{-1} F\bar{\Sigma} \end{bmatrix}$$
$$= \begin{bmatrix} 1-\sigma_1 s_{11} & -F\bar{\Sigma} \\ 0 & I - \bar{S}\bar{\Sigma} \end{bmatrix},$$

where $\bar{S} \equiv H - \sigma_1 G(1-\sigma_1 s_{11})^{-1} F$.

As $\bar{\Sigma}$ is factored out, all one needs to calculate $\bar{S}$ and $1-\sigma_1 s_{11}$ and thus be able to perform the first elimination step on the system matrix, is to choose a value for σ_1.

Moreover, if we denote by c' the updated vector $\hat{c}$ after one step of Gaussian elimination and define $\bar{c} \equiv (c'_2, \ldots, c'_n)^T \in \mathbb{R}^{n-1}$, then it is

$$\bar{c}_i \equiv \hat{c}_{i+1} - \frac{\sigma_1 \cdot (-s_{i+1,1})}{1 - \sigma_1 \cdot -s_{11}} \cdot \hat{c}_1 = \hat{c}_{i+1} + \frac{\sigma_1 \cdot s_{i+1,1}}{1 - \sigma_1 \cdot -s_{11}} \cdot \hat{c}_1 \tag{17}$$

for all $i \in [n-1]$. And again this transformation can be performed, once σ_1 is fixed. Hence, one step of Gaussian elimination can be performed on the system (4) if σ_1 is fixed.

For the proper value of σ_1 we get a correct step. Here by correct we mean that the unique solution $\bar{z}$ of the reduced system equals the vector $(z_2, \ldots, z_n)^T$, that is, the elimination step is correct if the solution of the reduced system is identical to the last $n-1$ components of the solution z of $(I - S\Sigma)z = \hat{c}$. One could, of course, also perform an elimination step with the wrong sign-choice. But then the equality of $\bar{z}$ and $(z_2, \ldots, z_n)^T$ would get lost. This shows:

Lemma 4.1. *If the sign of z_1 is known, then one correct step of Gaussian elimination can be performed on system* (4)*.*

Furthermore, for $\|S\|_\infty < 1$, the diagonal of the matrix $M \equiv I - S\Sigma$ is strictly positive. Since neither row and column pivots, nor multiplication with a signature [recall (5)] change the infinity norm of S, this shows:

Lemma 4.2. *If $\|S\|_\infty < 1$, no row or column pivot in S leads to numerical instabilities in the performance of a Gaussian elimination step on* (4)*.*

Thus, we can always, *by symmetric row and column pivoting* (which is also called *full pivoting* in some sources), produce a constellation for the AVE, where $\hat{c}_1 \in C_{\max}$. Then Theorem 3.1. provides us with the knowledge of the correct σ_1, if S conforms to any of the conditions listed in the main result. If we want to perform more than only the first step of a Gaussian elimination applying this principle, we need the constraints to hold for the reduced subsystem(s) as well. The following technical lemma ensures this for all structural restrictions stated in the main theorem.

Lemma 4.3. *Let $S \in \mathrm{M}_n(\mathbb{R})$ with $\|S\|_\infty = \xi < 1$, and define $\bar{S}$, $\bar{\Sigma}$ and E, F, G, H as in* (16)*. Then the following statements hold:*

1. $\|\bar{S}\|_\infty \leq \xi < 1$.
2. *If S is strictly diagonally dominant, then so is $\bar{S}$.*
3. *If S is symmetric, then so is $\bar{S}$.*

4. If S is tridiagonal, then so is $\bar{S}$.

The proofs can be looked up in the appendix of [Rad16]. Note that *4.* holds for arbitrary bandwidths of S.

Remark 4.4. *It is not hard to find structural restrictions that allow for a loosening of the norm constraints on S, while $C_{\max} \cap \Sigma_{\neq} = \emptyset$ holds. The difficulty is that these restrictions have to be invariant under the reduction steps of a Gaussian elimination, which consist of a mere addition of an outer product to the subsystem. For example, antisymmetry will necessarily get lost, as the only antisymmetric outer product is the zero matrix.*

4.2 The algorithm

The key idea for the signed Gaussian elimination is simple: Pivot the entry in $C_{\max}$ with the smallest index (and the corresponding row and column) to the first position, assume its sign to be correct and set it as the σ_1 for the first elimination step. Then repeat the procedure for the reduced system and so forth.

Below is a pseudocode for the algorithm. It makes use of the following conventions:

- $S, z, \hat{c}$ and Σ are defined as in (4).
- P_{jk} denotes the permutation matrix that corresponds to a transposition of j and k.
- $i^{jn}_{C_{\max}}$ denotes the smallest index i with $j \leq i \leq n$, where $\hat{c}_i \in C_{\max}$.
- *GaussStep*(A, b, j) is the signature of a function that performs the j-th step of a Gaussian elimination on a system, where $Ax = b$.

4.3 Correctness

With the results gathered so far, the proof of correctness for the conditions described in Theorem 3.1. is little more than a formality:

Proposition 4.5. *The SGE computes the unique solution of* (4) *correctly, if S conforms to any of the conditions described in Theorem 3.1.*

Algorithm Signed Gaussian elimination

$P = I$
for $j = 1 : n$ **do**
  $k = i^{jn}_{C_{\max}}$
  $\sigma_k = 1$
  if $\hat{c}_k < 0$ **then**
    $\sigma_k = -1$
  end if
  $S = P_{jk} S P_{jk}$
  $\Sigma = P_{jk} \Sigma P_{jk}$
  $\hat{c} = \hat{c} P_{jk}$
  $P = P P_{jk}$
  $(I - S\Sigma, \hat{c}) = GaussStep(I - S\Sigma, \hat{c}, j)$
end for
$z = (I - S\Sigma)^{-1} \hat{c}$
$z = zP$
return z

Proof. For all cases we have $\rho_0^s(S) \leq \|S\|_\infty < 1$, which guarantees the unique solvability of (4) and allows for unproblematic (symmetric) pivoting of rows and columns (Theorem 2.1. and Lemma 4.2.). Theorem 3.1. guarantees the correctness of the first sign choice. Lemma 4.3. assures that the conditions of the theorem are also satisfied by the reduced system. Hence, the argument applies recursively.

For the tridiagonal case we remark that for $n = 1$ and $\|S\|_\infty < 1$ we always have $\mathbf{sign}(z_1) = \mathbf{sign}(\hat{c}_1)$. Hence, the reduction step from a two- to a one-dimensional subsystem is unproblematic with regard to the correctness of the result. Even though for a square matrix of dimension one the notion of tridiagonality clearly makes no sense. □

The proposition shows that, in a way, systems which conform to the conditions of the main theorem behave like dented linear systems rather than fully fledged piecewise linear systems.

Remark 4.6. *Let S and z be generated uniformly at random. Then the expected value of $S|z|$ is the zero vector. This means, even though the infinity norm of S may be arbitrarily large, for $\hat{c} \equiv z - S|z|$ the sign of $\hat{c}_i$ is a maximum likelihood estimate for the sign of z_i for all $i \in [n]$ (n the dimension of the system). So, for the SGE any of the popular testing of algorithms beyond their proven correctness*

range with randomly generated systems would be a rather pointless exercise: Relevant problem dimensions begin in the thousands, where the law of large numbers makes a false estimate highly unlikely.

4.4 Effect on runtime

Throughout this analysis we will assume a uniform cost model, i.e., elementary arithmetic operations, as well as reading, writing and comparing a floating point number are all assumed to be in $\mathcal{O}(1)$. It is well known that, within this model, the Gaussian eliminations for dense and tridiagonal matrices have a complexity in $\mathcal{O}(n^3)$ and $\mathcal{O}(n)$, respectively. (See, e.g., [CLRS07, p. 752], and [CLRS07, p. 769].) The SGE has three types of additional operations in comparison to a classical Gaussian elimination *without* pivoting:

1. Determining the entry in $C_{\max}$ with the smallest index before every elimination step.

2. Pivoting in S and $\hat{c}$ before the elimination step.

3. Permuting the entries of the solution into their correct order after the completed backwards substitution.

For dense matrices this means that the SGE has precisely the cost of a Gaussian elimination with symmetric row/column-pivoting, which is roughly $\frac{1}{3}n^3$ fused multiply-adds (see the above references). So the SGE for dense matrices has the same asymptotical complexity as the unaltered algorithm. (For a detailed account of the operations of the different types of Gaussian elimination, see, e.g., [CLRS07, pp. 744-752].)

For tridiagonal systems the second and third point can clearly be handled in $\mathcal{O}(n)$. However, an analysis of the first point shows that the additional operations increase the asymptotical complexity of the tridiagonal SGE in comparison to the tridiagonal Gaussian elimination:

For simplicity we assume that every elimination step produces no zeros beyond the column that is eliminated. That is, the reduced subsystems stay densely tridiagonal in the sense that the three diagonals have no zero entries. Then in every column there is exactly one nonzero entry below the principal diagonal. Thus, the i-th elimination step exclusively affects row $(i+1)$ of S and thus only entry $(i+1)$ of $\hat{c}$. That is to say: $\hat{c}_{i+2}$ to $\hat{c}_n$ remain unaltered. Accordingly, it would be inefficient to run a comparison of all remaining entries of $\hat{c}$ after each elimination step. We outline a better approach:

Assume that $\hat{c}_{i+2}$ to $\hat{c}_n$ are sorted by absolute value (highest first) before the i-th elimination step. The only entry of $\hat{c}$ updated in the i-th step is $\hat{c}_{i+1}$. Then, to determine the entry between $\hat{c}_{i+1}$ and $\hat{c}_n$ with the largest absolute value, one only has to compare $\hat{c}_{i+1}$ and $\hat{c}_{i+2}$. The only entry of $\hat{c}$ updated in the next elimination step (after swapping $\hat{c}_{i+1}$ and $\hat{c}_{i+2}$, if necessary) is $\hat{c}_{i+2}$. And $\hat{c}_{i+3}$ to $\hat{c}_n$ remain sorted by absolute value. Hence, the argument applies recursively.

Now let $i = 1$, i.e., sort $\hat{c}$ before the first elimination step. Then, afterhand we need only $n-1$ comparisons and at most $n-1$ swaps throughout the elimination, which is clearly in $\mathcal{O}(n)$. As sorting n floats has the trivial lower bound $\mathcal{O}(n)$, but is currently not possible with this efficiency, the overall complexity of determining the proper order for the elimination is bounded from below by the complexity of the utilized sorting algorithm.

Asymptotically this approach is optimal, since determining the entry of $\hat{c}$ with the largest absolute value in every step clearly has sorting $\hat{c}$ once as a lower bound.

Hence, the tridiagonal SGE is at least as expensive as the algorithm utilized to sort $\hat{c}$. Since the tridiagonal Gaussian elimination's complexity is in $\mathcal{O}(n)$, this especially means that the tridiagonal SGE has – at the present state of research – a higher asymptotical complexity than the unmodified algorithm. (Note that, apart from a higher constant factor, this result holds for any fixed bandwidth of S.)

Currently, the asymptotically fastest sorting algorithm for floating point numbers, developed by Han and Thorup in [HT02], has a complexity of $\mathcal{O}(n \cdot \sqrt{\log \log n})$. But this is only a theoretical performance, since the latter is inefficient for realistic problem dimensions. For an actual application the use of an easily implementable in-place sorting algorithm such as Quicksort with its $\mathcal{O}(n \cdot \log n)$ average cost (see, e.g., [CLRS07, pp. 143-161]) is a far more adequate choice.

5 Sharpness of the bounds

For $n = 1$ we always have $\mathbf{sign}(z_1) = \mathbf{sign}(\hat{c}_1)$, if $\|S\|_\infty < 1$. So, naturally, we are inclined to ask whether the bounds from Theorem 3.1. can be loosened further.

Proposition 5.1. *Let $S \in \mathrm{M}_n(\mathbb{R})$ with $\|S\|_\infty \leq \frac{1}{2}$, and $z, \hat{c} \in \mathbb{R}$ s.t. $z + S|z| = \hat{c}$. Then for $n \geq 2$ the following holds:*

1. *It is possible that there exists a $c_i \in C_{\max}$ such that $\mathbf{sign}(\hat{c}_i) \neq \mathbf{sign}(z_i)$.*
2. *If $c_i \in C_{\max}$ and $\mathbf{sign}(\hat{c}_i) \neq \mathbf{sign}(z_i)$, then $z_i = 0$.*

Proof. 1.: Let

$$S \equiv \begin{bmatrix} 0 & \frac{1}{2} & 0 & . & 0 \\ 0 & \frac{1}{2} & 0 & . & 0 \\ . & . & . & . & . \\ 0 & \frac{1}{2} & 0 & . & 0 \\ 0 & \frac{1}{2} & 0 & . & 0 \end{bmatrix} \qquad \text{and} \qquad z \equiv \begin{bmatrix} 0 \\ 1 \\ \dots \\ 1 \\ 1 \end{bmatrix},$$

then $\hat{c} = (-\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})^T$ and we clearly have $c_1 \in C_{\max}$, but $\mathbf{sign}(\hat{c}_1) \neq \mathbf{sign}(z_1)$.

2.: We first exclude the following special case from the discussion: If we have $0 = \hat{c}_i \in C_{\max}$, then $\hat{c} = \mathbf{0}$ and thus, by unique solvability of the system, z is the zero vector as well. That is, we cannot have $\mathbf{sign}(\hat{c}_i) \neq \mathbf{sign}(z_i)$.

Now assume there were a $0 \neq \hat{c}_i \in C_{\max}$ with $\mathbf{sign}(z_i) \neq \mathbf{sign}(\hat{c}_i)$ and $z_i \neq 0$. Let $z_j \in Z_{\max}$. Since $\|S\|_\infty \leq \frac{1}{2}$, we have:

$$\frac{1}{2} \cdot |z_j| \geq \left|e_k^T S|z|\right| \qquad \forall k \in [n]. \tag{18}$$

As $\hat{c}_j = z_j - e_j^T S|z|$, this especially gives $|\hat{c}_j| \geq \frac{1}{2} \cdot |z_j|$. We proceed by a case distinction:

Case 1: Let $z_i \in Z_{\max}$. If z_i were zero, then z would be the zero vector and thus $\hat{c}$ as well. If we had $|z_i| > 0$, then $\hat{c}_i$ would have to adopt the sign of z_i due to (18). Hence, we would have $\mathbf{sign}(z_i) = \mathbf{sign}(\hat{c}_i)$ in contradiction of the initial assumption.

Case 2: Let $z_i \notin Z_{\max}$. Since $\mathbf{sign}(z_i) \neq \mathbf{sign}(\hat{c}_i)$, but $\hat{c}_i = z_i - e_i^T S|z|$, the sign of $\hat{c}_i$ must be the same as that of $-e_j^T S|z|$. This gives the leftmost inequality in

$$|\hat{c}_i| \leq \left|e_i^T S|z|\right| \leq \frac{1}{2} \cdot |z_j| \leq |\hat{c}_j|,$$

where $z_j \in Z_{\max}$. Since $\hat{c}_i \in C_{\max}$, we have $|\hat{c}_i| = |e_i^T S|z|| = |\hat{c}_j|$ – and the left equality clearly yields $z_i = 0$. □

The second statement of the proposition makes sure that the SGE calculates the proper solution for arbitrarily structured S with $\|S\|_\infty \leq \frac{1}{2}$, while the statement of the main theorem does not hold anymore in its strict sense that $z_i = 0$ if and only if $0 = \hat{c}_i \in C_{\max}$. That is, the SGE also computes solutions on orthant boundaries correctly.

Also note that, by replacing the inequalities in (18) with strict inequalities, the proof of Proposition 5.1.2. can be used as an alternate proof for Theorem 3.1.1.

One might ask now, if the SGE still runs provably correct with irreducible S that have a norm greater than one half. We will see below that the answer

to this query is no. Accordingly, under purely practical considerations the first two points of Theorem 3.1. could be merged into one condition: S arbitrarily structured with $\|S\|_\infty \leq \frac{1}{2}$.

Proposition 5.2. *For an irreducible $S \in \mathrm{M}_n(\mathbb{R})$ with $n \geq 2$, the correctness of the SGE cannot be ensured, if $\|S\|_\infty > \frac{1}{2}$.*

Proof. We start by demonstrating the sharpness of the bound for $n = 2$. For an $\epsilon > 0$ let

$$S \equiv \begin{bmatrix} \frac{\epsilon}{2} & \frac{1+\epsilon}{2} \\ 0 & \frac{1}{2} \end{bmatrix} \qquad \text{and} \qquad z \equiv \begin{bmatrix} \frac{\epsilon}{2} \\ 1 \end{bmatrix}.$$

Then, for $\hat{c} \equiv z - S|z|$ we have $\hat{c} = (-\frac{2+\epsilon}{4}, \frac{1}{2})^T$. And clearly $|c_1| > |c_2|$, but $\mathbf{sign}(\hat{c}_1) \neq \mathbf{sign}(z_1)$.

The structure of this example can be extended to higher dimensions. Let

$$S \equiv \begin{bmatrix} \frac{\epsilon}{2} & \frac{1+\epsilon}{2} & 0 & . & 0 \\ 0 & 0 & 0 & . & \frac{1}{2} \\ . & . & . & . & . \\ 0 & 0 & \frac{1}{2} & . & 0 \\ 0 & \frac{1}{2} & 0 & . & 0 \end{bmatrix} \qquad \text{and} \qquad z \equiv \begin{bmatrix} \frac{\epsilon}{2} \\ 1 \\ \dots \\ 1 \\ 1 \end{bmatrix}.$$

This yields $\hat{c} = (-\frac{2+\epsilon}{4}, \frac{1}{2}, \dots, \frac{1}{2})^T$. And again: $c_1 \in C_{\max}$, but $\mathbf{sign}(\hat{c}_1) \neq \mathbf{sign}(z_1)$.

As in both cases S is irreducible with $\|S\|_\infty = \frac{1}{2} + \epsilon$, this establishes the sharpness of the bound for $n \geq 2$. □

In the tridiagonal case the bound is sharp:

Proposition 5.3. *Let $S \in \mathrm{M}_n(\mathbb{R})$ be tridiagonal and symmetric. If $\|S\|_\infty \geq 1$, then the correctness of the SGE cannot be ensured.*

Proof. Just consider $S = -I$ and z the vector with entries -1. Then $\hat{c} = z + |z| = \mathbf{0}$ – and the SGE fails, since it picks $+1$ as σ_1. □

On a more general note, keep in mind that for $\|S\|_\infty \geq 1$ the unique solvability cannot be guaranteed anymore, which means that we enter an altogether different problem sphere.

We did not manage to find a counterexample that establishes the absolute sharpness of the bound in Theorem 3.1.3. However, we can demonstrate that the norm constraint for strictly diagonally dominant matrices can be loosened at most by a minute amount:

Proposition 5.4. *Let $S \in \mathrm{M}_n(\mathbb{R})$ be strictly diagonally dominant. If $\|S\|_\infty \geq \frac{2}{3} + \frac{1}{3(n+1)}$, then the correctness of the SGE cannot be ensured.*

Proof. Let

$$S \equiv \begin{bmatrix} \frac{1}{3} + \frac{1}{3(n+1)} & \frac{1}{3(n-1)} & \frac{1}{3(n-1)} & . & \frac{1}{3(n-1)} \\ 0 & \frac{2}{3} + \frac{1}{3(n+1)} & 0 & . & 0 \\ 0 & 0 & \frac{2}{3} + \frac{1}{3(n+1)} & . & 0 \\ . & . & . & . & . \\ 0 & 0 & 0 & . & \frac{2}{3} + \frac{1}{3(n+1)} \end{bmatrix} \quad \text{and} \quad z \equiv \begin{bmatrix} \frac{n+1}{2n+1} \cdot \epsilon \\ 1 \\ 1 \\ \dots \\ 1 \end{bmatrix}.$$

Then $\hat{c} = (-\frac{1}{3} \cdot (1-\epsilon), \frac{1}{3} \cdot \frac{n}{n+1}, \dots, \frac{1}{3} \cdot \frac{n}{n+1})^T$. Now choose an $\epsilon > 0$ such that $1 - \epsilon > \frac{n}{n+1}$. Then $\hat{c}_1$ has the largest absolute value of all entries in $\hat{c}$, but $\mathbf{sign}(z_1) \neq \mathbf{sign}(\hat{c}_1)$ – even in the strict sense that $\hat{c}_i < 0$, but $z_i > 0$. Hence, the first sign choice of the SGE fails. □

6 Acknowledgements

The author is indebted to his teacher Andreas Griewank and his colleague Tom Streubel, upon whose research the findings presented above are built. He also wants to convey his gratitude to Dorothee Schueth, Gestur Olafsson, and Jens Uwe Bernt for the time they invested into discussing and proofreading this work. A special thanks goes to Pascal Lenzner and Katharina Klost for their help with the presentation of the algorithmic section.

References

[BC08] L. Brugnano and V. Casulli. Iterative solution of piecewise linear systems. *SIAM Journal on Scientific Computing*, 30(1):463–472, 2008.

[CLRS07] T. H. Corman, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. Oldenbourg, 2007.

[CPS92] R. W. Cottle, J.-S. Pang, and R. E. Stone. *The Linear Complementarity Problem*. Academic Press, 1992.

[GBRS15] A. Griewank, J. U. Bernt, M. Radons, and T. Streubel. Solving piecewise linear equations in abs-normal form. *Linear Algebra and Its Applications*, 471:500–530, 2015.

[HHZ11] S.-L. Hu, Z.-H. Huang, and Q. Zhang. A generalized newton method for absolute value equations associated with second order cones. *Journal of Computational and Applied Mathematics*, 235(5):1490–1501, 2011.

[HT02] Y. Han and M. Thorup. Integer sorting in $\mathcal{O}(n\sqrt{\log\log n})$ and linear space. *Proceedings of the 43rd Symposium on Foundations of Computer Science.*, pages 135–144, 2002.

[JT95] C. R. Johnson and M. J. Tsatsomeros. Convex sets of nonsingular and p-matrices. *Linear and Multilinear Algebra*, 38(3):233–239, 1995.

[Kit98] B. Kitchens. *Symbolic dynamics: one-sided, two-sided and countable state markov shifts.* Springer, 1998.

[Man07a] O. .L. Mangasarian. Absolute value equation solution via concave minimization. *Optimization Letters*, 1(1):3–8, 2007.

[Man07b] O. .L. Mangasarian. Absolute value programming. *Computational Optimization and Applications*, 36(1):43–53, 2007.

[Man14] O. L. Mangasarian. Absolute value equation solution via linear programming. *Journal of Optimization Theory and Applications*, 161(3):870–876, 2014.

[MM06] O. .L. Mangasarian and R.R. Meyer. Absolute value equations. *Linear Algebra and Its Applications*, 419:359–367, 2006.

[Neu90] A. Neumaier. *Interval methods for systems of equations.* Cambridge University Press, 1990.

[Rad16] M. Radons. *Efficient solution of piecewise linear systems.* Master-Thesis, 2016.

[Roh89] J. Rohn. Systems of linear interval equations. *Linear Algebra and Its Applications*, 126:39–78, 1989.

[Rum97] S. M Rump. Theorems of perron-frobenius type for matrices without sign restrictions. *Linear Algebra and Its Applications*, 266:1–42, 1997.

[Sch12] S. Scholtes. *Introduction to piecewise differentiable equations.* Springer, 2012.

[YY12] X.-T. Yuan and S. Yan. Nondegenerate piecewise linear systems: A finite newton algorithm and applications in machine learning. *Neural Computation*, 24(4):1047–1084, 2012.